

Heterogeneous Knowledge Transfer in Video Emotion Recognition, Attribution and Summarization

Baohan Xu, Yanwei Fu, Yu-Gang Jiang, Boyang Li and Leonid Sigal

Abstract—Emotional content is a key element in user-generated videos. However, it is difficult to understand emotions conveyed in such videos due to the complex and unstructured nature of user-generated content and the sparsity of video frames that express emotion. In this paper, for the first time, we study the problem of transferring knowledge from heterogeneous external sources, including image and textual data, to facilitate three related tasks in video emotion understanding: emotion recognition, emotion attribution and emotion-oriented summarization. Specifically, our framework (1) learns a video encoding from an auxiliary emotional image dataset in order to improve supervised video emotion recognition, and (2) transfers knowledge from an auxiliary textual corpus for zero-shot recognition of emotion classes unseen during training. The proposed technique for knowledge transfer facilitates novel applications of emotion attribution and emotion-oriented summarization. A comprehensive set of experiments on multiple datasets demonstrate the effectiveness of our framework.

Index Terms—Video, Emotion Recognition, Transfer Learning, Zero-Shot Learning, Summarization.

1 INTRODUCTION

RAPID development of mobile devices has led to an explosive growth of user-generated images and videos, which creates a demand for computational understanding of visual media content. In addition to recognition of objective content, such as objects and scenes, an important dimension of video understanding is emotional or affective content. Such content can strongly resonate with viewers and plays a crucial role in the video-watching experience. Some successes have been achieved with the use of deep-learning architectures trained for text at both sentence- and document-level [39] or image sentiment analysis [8]. However, the ability to understand emotions from video, to a large extent, remains an unsolved problem.

The understanding of emotion information in videos has many real-world applications. Video recommendation services, such as those employed by YouTube and Netflix, can benefit from matching user interests with the emotions of video content and prediction of interestingness [23], [24], [35], leading to improved user satisfaction. Better understanding of video emotions may enable the placement of advertisements that are consistent with the main video’s mood and help avoid social inappropriateness such as placing a funny advertisement alongside a funeral video. Video summarization can also benefit from understanding emotions, since a good summary should probably include strong emotional content from the original video.

This paper tackles three inter-related problems in video emotion understanding. We start with emotion recogni-

tion in both supervised and zero-shot conditions. Zero-shot video emotion recognition aims to recognize emotion classes that are not seen during training. This task is motivated by recent cognitive theories [3], [4], [7], [44] that suggest human emotional experiences extend beyond the traditional “basic emotion” categories (e.g. Ekman’s six emotions [16]). Emotion processes and other cognitive processes in our brain cooperate closely to create rich and diverse emotional and affective experiences [27], [43], [52], such as ecstasy, nostalgia, or suspense. Therefore, when operating in the real world, recognition systems trained with a small set of emotion labels will inevitably encounter emotional and affective expressions not in its training set. In zero-shot learning, we rely on heterogeneous sources of background knowledge to alleviate reliance on training data and improve generalization to labels previously unseen. Since our goal is to recognize both emotional and affective impact of video from a computer vision perspective, we use “emotion” as a shorthand to refer to both emotion and affect. the boundary between which can be blurry (see, for example, the argument around surprise [57].)

We then define a novel task called *video emotion attribution*, which aims to identify the contribution of each frame in a video to its overall emotion. Not every frame contributes emotional information. This realization further enables the third task – emotion-oriented video summarization. In the emotion-oriented video summarization, we summarize a video and at the same time preserve emotional content as much as possible.

Our knowledge transfer framework performs two types of knowledge transfer. In supervised emotion recognition, we learn an effective video encoding from a large-scale auxiliary emotional image dataset. This auxiliary Image Transfer Encoding (ITE) process can generate a video representation more conducive to emotion recognition than al-

- Baohan Xu and Yu-Gang Jiang are with the School of Computer Science, Shanghai Key Lab of Intelligent Information Processing, Fudan University, Shanghai, China. Email: {bhxu14, ygj}@fudan.edu.cn.
- Yanwei Fu, Boyang Li and Leonid Sigal are with Disney Research Pittsburgh, PA, USA. Email: y.fu@qmul.ac.uk, {lsigal, albert.li}@disneyresearch.com.



Fig. 1: **Humans make use of contextual information in emotion recognition.** The face of U.S. Senator Jim Webb, when inspected in isolation on the left, appears agitated and angry. However, when put back in the context on the right, he appears happy and excited. Image credit: Doug Mills/The New York Times. Reproduced from [4].

ternative methods. In zero-shot emotion recognition, where the emotion to be recognized is not in the training set, we learn embeddings for emotional words from a large-scale text corpus, so that an unknown emotion word can be semantically related to labeled or known emotions and subsequently recognized. The knowledge transfer technique further facilitates two novel applications: emotion attribution and emotion-oriented summarization. Video sub-shots more related to the overall emotion are identified and emphasized in the automatically generated summary.

Contributions: We introduce a framework for transferring knowledge from heterogeneous sources (image and text) for the understanding of emotions in videos. For the first time, we demonstrate zero-shot emotion recognition by utilizing knowledge learned from text sources to the video domain. We also propose the first definitions and solutions for the problems of emotion video attribution and emotion-oriented summarization. We show that our emotion-oriented summaries are better than alternative methods that do not consider emotion. Finally, we introduce and will make available to the community two new emotion-centric video datasets: the Plutchik and Ekman emotion datasets.

2 RELATED WORK

2.1 Psychological Theories of Emotion

It is a widely held view in psychology that emotion contains a number of static categories, each of which is associated with stereotypical facial expressions, physiological measurements, behaviors, and external causes [15], [17]. The most well known model is probably Ekman’s six pan-cultural basic emotions, including happiness, sadness, disgust, anger, fear, and surprise [16], [17]. However, the exact categories can vary from one model to another. Plutchik [59] added anticipation and trust to the list. Ortony, Clore and Collins’s [56] model of emotion defined up to 22 emotions, including categories like hope, shame, and gratitude.

Nevertheless, more recent empirical findings and theories from the psychological construction approach [3], [44] suggest emotional experiences are much more varied than previously assumed. It is argued that the categories are modal or stereotypical emotions, and large fuzzy areas exist

in the emotion landscape. Instead of associating a fixed set of facial expressions with each emotion, emotion recognition is influenced by visual and language context [4], [7]. The facial expression of smile, for example, can indicate happiness, embarrassment, or being subordinate in different contexts [3]. See Fig. 1 for a real-world example.

Other theories [27], [43], [52] highlight the dynamics of emotion and the interactions between emotional processes and other cognitive processes. Together, the complex dynamics and interactions lead to a rich set of emotional and affective experiences, and correspondingly rich natural language descriptions, such as ecstasy, nostalgia, or suspense. In order to cope with diverse emotional descriptions that may be practically difficult (or at least very costly to label), in this paper we investigate emotion recognition in a zero-shot setting (in addition to, a traditional, supervised setting). Our emotion recognizer is tested against emotional classes that do not appear in the training set. The zero-shot recognition task is designed to test the system’s ability to make use of knowledge learned from heterogeneous sources in order to adapt to unseen tags.

2.2 Automatic Emotion Analysis

In this section, we briefly review three relevant areas of research: recognition of facial expressions from images and videos, recognition of emotional impact of images on viewers, and recognition of emotional impact from videos.

Recognition of facial expressions from images and videos.

In the light of findings that humans rely on contextual information to recognize emotion [4], [7], in this paper, we aim to recognize the overall emotional impact of a video from all of its information. This contrasts with the identification of facial expressions and the associated emotion in static images and video, which has been a subject extensively studied. Two recent reviews of the topic can be found in [58] and [65]. Several competitions, such as the Facial Expression Recognition and Analysis Challenge [67], the Audio/Visual Emotion Challenge [60], and the Emotion Recognition In The Wild Challenge [12], have been held. Notably, Liu *et al.* [47] construct a mid-level representation called expressionlet from spatio-temporal manifold. Cruz *et al.* [10] proposed a dynamic downsampling of facial expressions in video

at a rate proportional to the rate of temporal changes of visual information. When apex labels are provided, videos are downsampled around the apex of an expression.

A few recent works focused on predicting emotions, affects, and emotion-related cognitive states outside the basic emotion categories. Kaliouby and Robinson [18] recognized emotion-related mental states, such as agreeing or feeling unsure. Bosch *et al.* [6] detect learning-related affects including boredom, confusion, delight, engagement, and frustration; other work recognized smirk [62] / fatigue [31].

Recognizing the emotional impact of still images on viewers. Machajdik and Hanbury [51] classified images into 8 affective categories: amusement, awe, contentment, excitement, anger, disgust, fear, and sadness. In addition to color, texture, and statistics about faces and skin area present in the image, they also make use of composition features such as the rule of the third and depth of field. Lu *et al.* [48] studied shape features along the dimensions of rounded-angular and simple-complex, and their effects in arousing viewers' emotions. You *et al.* [77] designed a deep convolutional neural network (CNN) for visual sentiment analysis. After training on the entire training set, images on which the CNN performs poorly are stochastically removed. The remaining images were used to fine-tune the network. A few work [8], [74] also employed off-the-shelf CNN features.

Recognizing emotional impact from videos. For a more comprehensive review, we refer reader to the latest survey [71]. A large number of early work studied emotion in movies (e.g. [30], [37], [68]). Wang and Cheong [68] used an SVM with diverse audio-visual features to classify 2040 scenes in 36 Hollywood movies into 7 emotions. Jou *et al.* [36] worked on animated GIF files. Irie *et al.* [30] use Latent Dirichlet Allocation to extract audio-visual topics as mid-level features, which are combined with a Hidden-Markov-like dynamic model.

SentiBank [5] contains a set of 1,553 adjective-noun pairs, such as "beautiful flowers" and "sad eyes", and images exemplifying each pair. One linear SVM detector was trained for each pair. The best-performing 1,200 detectors provide a mid-level representation for emotion recognition. Chen *et al.* [8] replaced the SVM detectors with deep convolutional neural networks. Jiang *et al.* [33] explored a large set of features and confirmed the effectiveness of mid-level representations like SentiBank. In this work, we transfer knowledge learned from the same set of Flickr images for the purpose of video emotion analysis.

The implicit approach for recognizing the emotional impact of a video is to recognize emotions exhibited by viewers of that video. This clever trick delegates the complex task of video understanding to human viewers, thereby simplifying the problem. McDuff *et al.* [53] analyzed facial expressions exhibited by viewers of video advertisements recorded with webcams. Histogram of Oriented Gradient (HOG) features were extracted based on 22 key points on the faces. Purchase intent is predicted based on the entire emotion trajectory over time. Kapoor *et al.* [38] used video, skin conductance, and pressure sensors on the chair and the mouse to predict frustration when a user interacted with an intelligent tutoring system. However, the success of this

approach depends on the availability of a large number of human participants.

Nevertheless, all previous work are limited as they aim to predict emotion or sentiment classes present in the training set. In this work, we utilize knowledge from other domains like images and text in order to identify emotion classes unseen in the training set. In addition, we also investigate related practical applications like emotion-oriented video attribution and summarization.

2.3 Multi-Instance Learning

The knowledge transfer approach adopted in this work is related to multi-instance learning (MIL), which has been extensively studied in the machine learning community and utilized in other domains. We therefore briefly review related techniques in the following. MIL refers to recognition problems where each label is associated with a bag of instances, such as a bag of video frames. It has been used in many problems, such as drug activity prediction [14], speech recognition [61], image retrieval and classification [78]. The problem investigated in this work is intrinsically a multi-instance learning case as each video consists of many frame instances with possibly different emotions.

There are basically two branches of MIL algorithms. In the first branch, many works attempted to enable single-instance supervised learning algorithms to be directly applicable to multi-instance feature bags. This branch includes most of the early works on MIL [1], [63] such as miSVM [2], MIBoosting [75], Citation-kNN [69], MI-Kernel [25], among others. These algorithms achieved satisfactory accuracies in several applications, but most of them can only handle small or moderate-sized data. In other words, they are computationally expensive and cannot be applied to deal with large-scale video data.

The second branch of works are more recent, where researchers tried to solve the MIL problems by adapting multi-instance bags to specific points in the original instance space. Popular algorithms include constructive clustering based ensemble (CCE) [78], multi-instance learning based on the Fisher Vector representation (Mi-FV) [73] and multi-instance learning via embedded instance selection [9]. Inspired by these works, we encode the video frame bags into single-instance representations of the bag-level information. Our approach is different from this category of MIL algorithms in the following: (1) our emotion recognition task is a multi-class multi-instance problem, while most of the previous MIL algorithms aimed at binary classification; (2) We perform the encoding process by using auxiliary data like images, and demonstrate that transferring such knowledge is important for video emotion analysis.

2.4 Video Summarization

Video summarization has been studied for more than two decades. A complete review is beyond the scope of this paper and we refer readers to [66].

There are two main types of video summaries: key-frames [11], [19], [29], [45] and video skims [20], [55], [70], [72]. Video summarization has been explored for various types of content, including professional videos (e.g., movies or news reports) [55], [70], [72], surveillance videos [19],

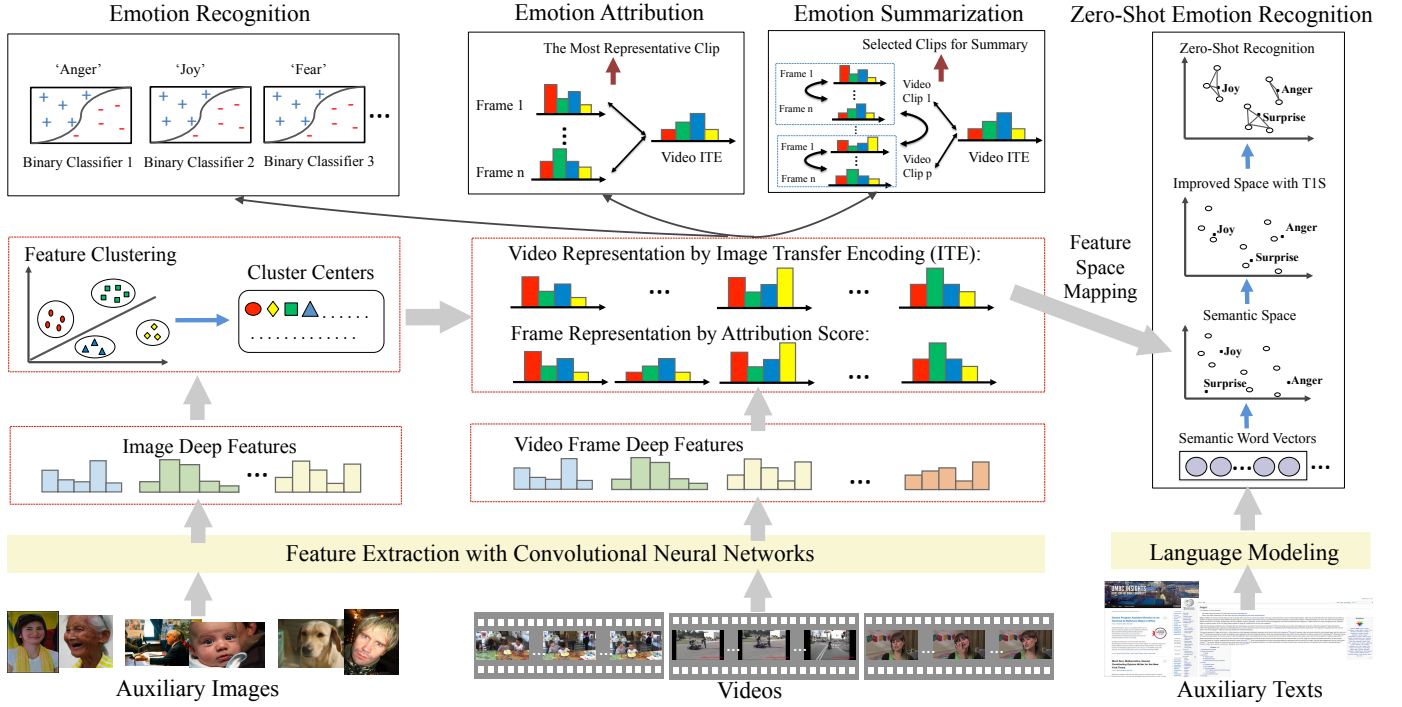


Fig. 2: **An overview of our framework.** Information from the auxiliary images (bottom left) is used to extract an emotion-centric dictionary from CNN-encoded image elements, which is subsequently used to encode video (bottom middle) and recognize emotion (top left). The same encoding is used for emotion attribution and summarization (top middle). Finally, information from a large text corpus is utilized for zero-shot recognition of emotions, as illustrated on the right.

[20], and, to a lesser extent, user-generated videos [72]. To extract the video summary, most approaches have to rely either on the low-level information, such as visual saliency [50] and motion cues [55]; or on the mid-level information e.g. object trajectories [45], tag localization [70] and semantic recognition [72]. Facial expressions has been considered by Dhall and Roland [13] who extracted video summaries by considering smile/happy face expression. However, none of these approaches have considered video summarization based on more general video emotion content. Such video emotion is an important cue for finding the most “interesting” or “important” video highlights. For example, a good summary of a birthday party, or a graduation ceremony, should capture the most emotional moments in the event. Not considering the valuable emotion dimension in the video summarization task risks losing these precious moments in the summary.

3 APPROACH

The overview of the proposed framework is illustrated in Figure 2. We start by introducing the problem formulation and common notation, and then discuss the auxiliary image transfer encoding for supervised recognition, auxiliary text-based transfer encoding for zero-shot recognition, and video emotion attribution and summarization.

3.1 Problem Setup

Suppose we have a training video dataset

$$Tr = \{(V_i, X_i, \mathbf{s}_i, z_i)\}_{i=1, \dots, n_{Tr}},$$

where the i^{th} video $V_i = \{\mathbf{f}_{i,1}, \dots, \mathbf{f}_{i,n_i}\}$ has n_i frames, with the features $X_i = \{\mathbf{x}_{i,1}, \dots, \mathbf{x}_{i,n_i}\}$, where the subscripts i, j denote the j^{th} frame of video V_i . $\mathbf{x}_{i,j}$ is extracted by the state-of-the-art deep Convolutional Neural Network (CNN) architecture which was recently shown to greatly outperform more traditional hand-crafted low-level features, such as HOG and SIFT, on several benchmark datasets including MNIST and ImageNet [40] in machine learning and computer vision communities. Specifically, we retrain AlexNet [40] with 2,600 ImageNet classes and use the seventh layer (“fc7”) representation as $\mathbf{x}_{i,j}$, computed on the input frame $\mathbf{f}_{i,j}$.

We use \mathbf{s}_i to denote the encoded video-level feature of video V_i obtained using auxiliary image transfer encoding, which is introduced in Section 3.2; $z_i \in Z_{Tr}$ is the class label of video V_i from the set of training labels Z_{Tr} and n_{Tr} is the total number of training videos. The testing data set is similarly defined as

$$Te = \{(V_i, X_i, \mathbf{s}_i, \tilde{z}_i)\}_{i=1, \dots, n_{Te}},$$

where n_{Te} is the total number of testing videos. Under the zero-shot learning setting, the labels in training and testing sets have different domains: $z_i \in Z_{Tr}$ and $\tilde{z}_i \in Z_{Tr} \cup Z_{Te}$, $Z_{Tr} \cap Z_{Te} = \emptyset$, denoting the testing set contains classes previously unseen.

To enable knowledge transfer, we introduce a large-scale auxiliary image set and a text sentiment dataset. We denote the auxiliary image sentiment dataset as $A = \{(a_i, \phi_i)\}_{i=1, \dots, |A|}$, where ϕ_i is the deep CNN feature of an image a_i which is extracted with the same trained model as $\mathbf{x}_{i,j}$ above.

The textual data are represented as a sequence of words $W = (w_0, \dots, w_{|W|})$, $w_j \in \mathcal{V}$ where the vocabulary \mathcal{V} is the set of unique words. We learn a \mathcal{K} -dimensional embedding ψ_w for each $w \in \mathcal{V}$.

3.2 Auxiliary Image Transfer Encoding (ITE) for Supervised Emotion Recognition

In multi-instance learning, the natural way of encoding multi-instance bags into single-instance representations is to cluster the instances of all the bags to several groups. Examples of this include CCE [78] and Mi-FV [73]. Such clustering enables the re-representations of each bag as the new Bag-of-Words (BoW) features. However, this comes with three challenges. First, the feature set is learned for the tasks of image classification, rather than for video emotion recognition. Second, the large and complex space of video visual elements (such as objects, events, and people) and their possible long-term temporal progression and interaction makes the domain intrinsically more complex than the previous image sentiment and MIL works. Third, the emotion is often expressed on certain limited (sparse) keyframes or video clips, and many video frames may not be directly related to the expressed emotion.

To address these challenges, we utilize emotion information from a large-scale emotional image dataset to encode the video content using a BoW representation. This can be intuitively explained from the perspective of entropy. A dictionary built from the auxiliary emotion-related images can efficiently encode a video frame with emotion information as a sparse vector which concentrates on a few dimensions. In comparison, a frame without emotion information will likely be encoded less efficiently, producing a denser vector with small values in many dimensions. As a result, a non-emotional frame will have higher entropy than the emotional frame, and hence less impact on the resulted BoW representation.

More formally, we first find D clusters from the auxiliary images and encode the video frames based on the cluster centers. Specifically, we perform a spherical k-means clustering on the auxiliary image dataset, which is to solve:

$$\min \sum_{i=1}^{|A|} (1 - \gamma_{i,d} \cos(\phi_i, \phi_d)). \quad (1)$$

The goal is to find D spherical cluster centers $\phi_1 \dots, \phi_D$. The cosine similarity is defined as

$$\cos(\phi_i, \phi_d) = \frac{\phi_i^\top \phi_d}{\|\phi_i\| \|\phi_d\|}. \quad (2)$$

The variable $\gamma_{i,d}$ assigns an image a_i to the closet cluster center d , which is defined as

$$\gamma_{i,d} = \begin{cases} 1 & \text{if } d = \operatorname{argmax}_j \cos(\phi_i, \phi_j), \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

To encode one video bag into its corresponding single-instance representations, a BoW scheme is used to translate the feature set X_i into a D -dimensional vector $\mathbf{s}_i = (s_{i,1}, \dots, s_{i,d}, \dots, s_{i,D})$. Specifically, to encode the feature vector of a video V_i , we fix the cluster centers $\{\phi_1, \dots, \phi_D\}$ found by the k-means and identify the K nearest cluster



Fig. 3: A frame from a graduation video that contains both smiling and weeping.

centers for each frame $\mathbf{f}_{i,j}$. We thus can get the assignments of each frame to each cluster, $\nu_{i,j,d}$:

$$\nu_{i,j,d} = \begin{cases} 1 & \text{if } \phi_d \in \text{KNN}(\mathbf{x}_{i,j}), \\ 0 & \text{otherwise,} \end{cases} \quad (4)$$

where $\text{KNN}(\mathbf{x}_{i,j})$ denotes the spherical K nearest neighbours¹ to $\mathbf{x}_{i,j}$ from the cluster centers $\{\phi_1, \dots, \phi_D\}$. We then accumulate the effects of each frame on the d^{th} dimension to compute the feature vector \mathbf{s}_i :

$$s_{i,d} = \sum_{j=1}^{n_i} \nu_{i,j,d} \cdot \cos(\mathbf{x}_{i,j}, \phi_d). \quad (5)$$

Our encoding scheme in Eq (5) is different from the standard BoW [64] and soft-weighting BoW encoding [34]. First, the traditional BoW encodes local descriptors, such as SIFT and STIP, which requires a dictionary orders of magnitude greater than our frame set. Thus directly using standard BoW [64] to our problem will make the generated video-level features too sparse to be discriminative. Second, the soft-weighting encoding assigns one visual feature point to multiple clusters using typically exponential or Gaussian kernel, to downweight the contribution to the clusters that are not closest to the feature. However, in our problem one video frame can express multiple emotions simultaneously. For example, Figure 3 shows an example of both grief and happiness. Thus in Eq (5), one feature instance $\mathbf{x}_{i,j}$ can equally contribute to different encoding bins. In other words, we make use of a uniform kernel instead.

The encoding scheme from the frame-level deep features to the video-level emotional representation helps the standard video emotion recognition tasks. Given a test video $V_k \in Te$, its class label can be estimated as

$$\hat{z}_k = \operatorname{argmax}_z \mathcal{L}(\mathbf{s}_k | \mathbf{s}_{Tr}, z_{Tr}), \quad (6)$$

where $\mathcal{L}(\cdot)$ is the predictor trained from the video-level feature set S_{Tr} of the training video set Tr . In this paper, we use the support vector machines (SVM) classifier with chi-square kernel as the predictor $\mathcal{L}(\cdot)$. The procedure of the ITE based supervised emotion recognition is summarized in Algorithm 1.

3.3 Zero-Shot Emotion Recognition

Canonical emotion theories such as Ekman [17] often provide detailed textual definitions for a fixed number of

1. Generally, we require that $K \gg 1$, since videos can express much more ‘versatile’ emotions than those of images (e.g. Figure 3).

Algorithm 1 Pseudo-code describing of ITE based supervised emotion recognition.

Require:

- A : an auxiliary image dataset.
 - Tr : training video set;
 - Te : testing video set.
- 1: Generate $\{\mathbf{x}_d\}$ and $\{\gamma_{a_i,d}\}$ by the spherical k-means on $A \leftarrow \text{Eq (1)}$;
 - 2: Generate a video-level feature set S_{Tr} and S_{Te} by encoding the videos in Tr and $Te \leftarrow \text{Eq (5)}$;
 - 3: Train chi-square SVM predictor $\mathcal{L}(\cdot)$ using S_{Tr} ;
 - 4: Recognize emotions in Te with the predictor $\leftarrow \text{Eq (6)}$.
-

prototypical emotions. However, recent research [3], [44] questioned the validity of basic emotional categories and highlighted differences within each category. This raises an interesting question: if we face a more diverse list of emotions than those in the training set, can we identify these emotions purely based on their textual description? This is the zero-shot recognition problem. To address it, we need to employ an auxiliary set of textual information to help encode the emotion classes that have never been visually seen before. Here, we use w_{Tr} and w_{Te} to indicate the emotion label words of auxiliary and testing dataset.

The encoding algorithm finds a distributed representation for each word in the vocabulary, by training models from large-scale textual corpora containing sentiment data. Specifically, we find a low-dimensional vector representation $\psi_w \in \mathbb{R}^K$ by using each word w_t in a corpus to predict nearby words by maximizing the following log probability:

$$\max J = \max \frac{1}{T} \sum_{t=1}^T \sum_{-c \leq j \leq c, j \neq 0} \log p(w_{t+j} | w_t), \quad (7)$$

$$p(w_o | w_i) = \frac{1}{Z} \exp(\psi_{w_o}^\top \psi_{w_i}),$$

where c is the context window for prediction; w_o and w_i are “input” and “output” words respectively, and ψ is the continuous word representation to learn. Optimizing J is an effective way to learn continuous word representations, but the total computational cost has been intractable until recent deep learning developments such as word2vec [54], and GloVe [26].

This acquired vector representation serves as the intermediary between an emotion word w and its corresponding video emotion class. Mapping video-level features into this semantic space requires a regressor from the video feature space to the word embedding space:

$$g : S_{Tr} \rightarrow \psi_{w_{Tr}}, \quad (8)$$

In this work, we train a support vector regressor with a linear kernel for each dimension of the word vector $\psi_{w_{Tr}}$. Similar support vector classifiers have also been used in the attribute learning works [22], [42].

Note that the regression models in Eq (8) potentially have a generalization problem which is largely caused by the different visual distribution of disjointed training and testing classes in zero-shot learning settings. For example, videos of joy usually have positive frames, whilst a sad video would have negative ones. To ameliorate such generalization problems, we take inspiration from [21] and

apply Transductive 1-Step Self-Training (T1S) to adjust the word vector of new emotion classes. Specifically, for a class $z^* \in Z_{Te}$ that is previously unseen, and the corresponding word vector ψ_{z^*} , we compute a smoothed version $\bar{\psi}_{z^*}$:

$$\bar{\psi}_{z^*} = \frac{1}{K} \sum_{g(\mathbf{s}_i) \in \text{KNN}(\psi_{z^*}), V_i \in Te} g(\mathbf{s}_i), \quad (9)$$

where $\text{KNN}(\psi_{z^*})$ denotes the set of spherical K nearest neighbors to ψ_{z^*} of the predicted testing word vector in the semantic space. Eq (9) aims at transductively ameliorating such visual differences by averaging the $\bar{\psi}_{z^*}$ with its nearest neighbor testing instances. Here, to prevent the semantic drift of self-training, we only do self-training for one step. By using Eq (9), we can get the updated word vector testing set $\bar{\psi}_{w_{Te}}$.

Thus, given a test video V_j in the testing set, its class label \tilde{z}_j can be estimated as

$$\tilde{z}_j = \underset{z \in Z_{Te}}{\operatorname{argmax}} \cos(g(\mathbf{s}_j), \bar{\psi}_z). \quad (10)$$

Compared with the zero-shot learning algorithm in [21], we skip the intermediate level of latent attributes and directly apply the 1-step self-training in the semantic word vector space. In addition, we use cosine similarity as the metric rather than the Euclidean distance. The process is summarized in Algorithm 2.

Algorithm 2 Pseudo-code describing the T1S algorithm.

Require:

- C : an auxiliary text dataset .
 - Tr : training video set;
 - Te : testing video set.
 - $Z_{Tr} \cap Z_{Te} = \emptyset$.
- 1: Train the word2vec language model with the large-scale auxiliary text dataset $\leftarrow \text{Eq (7)}$
 - 2: Project emotion words of training and testing sets to continuous vectors $\mathbf{v}_{w_{Tr}}$ and $\mathbf{v}_{w_{Te}}$;
 - 3: Train regressors $\leftarrow \text{Eq (8)}$;
 - 4: Perform zero-shot emotion recognition $\leftarrow \text{Eq (9) and Eq (10)}$.
-

3.4 Attribution and Summarization

We define the *video attribution* problem as one of attributing the emotion of a video to its frames/clips. The video emotion attribution problem is inspired by, and yet different from, text-based sentiment attribution [39]. The difference is that sentiment attribution only considers positive or negative attitudes while we consider more diverse emotions for emotion attribution.

Emotion attribution can help us find *video highlights* [29], which are the interesting or important events happened in the video. Generally, the concepts of “interesting” and “important” may be variable for different target video domains and applications, such as the scoring of a goal in soccer videos, applause and cheering in talk-show videos, and exciting speech in presentation videos. Nevertheless, most of these “interesting” or “important” video events may contribute to/convey very strong video emotions which are thus highlighting the core parts of the whole video.

Formally, for one video sequence $V_i = \{\mathbf{f}_{i,1}, \dots, \mathbf{f}_{i,n_i}\}$, we want to find the video frames that highly contribute to the overall expressed video emotion. Using Eq (2), we can encode a frame $\mathbf{f}_{i,j}$ as a similarity score vector:

$$\mathbf{h}_{i,j} = [\dots, \nu_{i,j,d} \cdot \cos(\mathbf{x}_{i,j}, \boldsymbol{\phi}_d), \dots], \quad (11)$$

where $\nu_{i,j,d}$ is defined in Eq (4). The vector $\mathbf{h}_{i,j}$ uses the auxiliary image dataset A to evaluate the emotions in the j^{th} frame.

Equipped with the frame-level emotion score vector, the video emotion attribution can be formulated as measuring the similarity between the video-level emotion vector and the frame-level vectors. Specifically, the attribution score of the j^{th} video frame is computed as the cosine similarity between \mathbf{s}_i , the feature vector of the entire video, and $\mathbf{h}_{i,j}$. To find the frame j^* that contributes the most to the overall emotion of V_i , we simply take the argmax over all frames:

$$j^* = \underset{j \in [1, \dots, n_i]}{\operatorname{argmax}} \cos(\mathbf{s}_i, \mathbf{h}_{i,j}). \quad (12)$$

Similarly, emotion attribution may be performed on a list of pre-partitioned video clips. Given a list of clips $\{E_1, \dots, E_P\}$, we can compute the similarity score vector for each clip as

$$\mathbf{h}_p = \left[\dots, \sum_{j \in E_p} \nu_{i,j,d} \cdot \cos(\mathbf{x}_{i,j}, \boldsymbol{\phi}_d), \dots \right], \quad (13)$$

and then perform a similar argmax operation.

Emotion attribution can facilitate interesting applications, such as video summarization and retrieval. To further enable the emotion-based *video summarization*, we can employ the emotion attribution technique described above and select a set of the most representative frames (or, the corresponding short clips containing the frames). Suppose VK_i is the set of key frames of the i^{th} video, defined as

$$VK_i = \underset{j \in VK_i}{\operatorname{argmax}} \cos(\mathbf{s}_i, \mathbf{h}_{i,j}) + \lambda \sum_{j \in VK_i} \sum_{k \in [1, \dots, n_i]} \cos(\mathbf{x}_{i,j}, \mathbf{x}_{i,k}), \quad (14)$$

where λ is a weight parameter, and the second term $\cos(\mathbf{x}_{i,j}, \mathbf{x}_{i,k})$ rewards key frames that are the most similar to other frames in the same video, which means that the selected set of key frames are representatives of the entire video.

Comparing with previous work [19], [20], [29], [66], [72], Eq (14) considers the summary of both video highlights (by the first term for emotion attribution) and information coverage (by the second term for eliminating redundancy and selecting information-centric frames/clips). Thus our method can produce a condensed, succinct and emotion-rich summary which can facilitate the browsing, retrieval and storage of the original video content. Particularly, our summary results are more emotional interpretable due to the emotion attribution.

4 EXPERIMENTS

4.1 Datasets and Settings

We adopt three video emotion datasets for evaluation. Among them, two (the Plutchik and the Ekman datasets)

are introduced by us and will be made available to the community once this paper is accepted.

The YouTube emotion dataset [33]. This dataset contains 1,101 videos annotated with 8 basic emotions from the Plutchik’s Wheel of Emotions. To better evaluate the zero-shot emotion recognition tasks, we re-annotate the videos into 24 emotions according to Plutchik’s definitions by adding 3 variations to each of the basic emotions. For example, we split the basic *joy* class into *ecstasy*, *joy*, and *serenity* along the arousal dimension. We use the shorthand **YouTube-8** and **YouTube-24** for the original and re-annotated datasets respectively. **YouTube-24** is used as a more difficult task for evaluating the zero-shot recognition.

The Plutchik dataset. This dataset is derived from the recently proposed VideoStory dataset [28]. We use the keywords of the Plutchik’s Wheel of Emotions [59] to query the textual descriptions of the VideoStory dataset. Since most of the textual descriptions come from the video captions, the emotions of the returned videos are accurately described by their corresponding emotional keywords. We manually pruned noisy videos from the returned set, which leads to a set of 626 videos belonging to 14 emotion classes.

The Ekman dataset. As discussed in the related work, the studies of Ekman [17] found a high agreement of emotions across cultures can be labelled as 6 basic emotion types. These 6 emotions are a subset of the 8 classes defined by Plutchik. Based on this, we collect the Ekman emotion dataset which has 1,637 videos in the 6 emotion classes, with a minimum number of 221 videos per class. These videos are collected from social video-sharing websites, such as YouTube and Flickr. The dataset is crowdsource labelled by 10 different annotators, each video was labeled by at least 3 of those annotators. Final annotation for each video is produced by a majority vote among the labels assigned by the annotators.

Auxiliary emotional image and text datasets. From the Flickr image dataset [5], we select as the auxiliary image data a subset of 110K images of Adjective-Noun Pairs (ANPs) that have top ranks with respect to the emotions (see Table 2 in [5]). These images are clustered into 2,000 clusters whose centers are used in Eq (2). As shown in [54], the large-scale text data can greatly benefit the trained language model. We train the Skip-gram model (Eq 7) on a large-scale text corpus, which includes around 7 billion words from the UMBC WebBase (3 billion words), the latest Wikipedia articles (3 billion words) and some other documents (1 billion words). The trained model has around 4 million unique words and phrases. Most of the documents are about scientific articles and professional reports which have very strict definitions, descriptions and usage of the emotion and sentiment related words.

Experimental settings. Each video is uniformly sampled at 5 frame increments for feature extraction to reduce the computational cost. The dimension of the real-valued semantic vectors $\boldsymbol{\psi}$ (Eq (7)) is set to 500. Our AlexNet CNN model is trained by ourselves using 2,600 ImageNet classes with the Caffe toolkit [32], and we use the 4,096-dimensional activations of the 7th fully-connected layer after the Rectified

Linear Units as features. The number of nearest neighbors in Eq (4) is empirically set to 10% of the image clusters, which balances the computational cost with a good representation in Eq (5).

4.2 Video Emotion Recognition

4.2.1 Supervised Recognition

We first report results on the supervised emotion recognition task, and compare our ITE encoding method with the following alternative baselines.

MaxP [46]. The instance-level classifiers are trained using the labels inherited from their corresponding bags. These classifiers can be used to predict instance labels of testing videos. The final bag labels are produced by majority vote of instance labels. This method is a variant of the Key Instance Detection (KID) [46] in multi-class multi-instance setting.

AvgP [76]. We average the frame-level image features of one video as video-level feature descriptions for classification. For the i^{th} video, its average pooling feature is computed as $\frac{1}{n_i} \sum_{j=1}^{n_i} \mathbf{x}_{i,j}$. The average pooling is the standard approach of aggregating frame-level features into video-level descriptions as mentioned in [76].

Mi-FV [73]. MIL bags of training videos are mapped into a new bag-level Fisher Vector representation. Mi-FV is able to handle large-scale MIL data efficiently.

CCE [78]. The instances of all training bags are clustered into b groups, and each bag is re-represented by b binary features, where the values of the i^{th} feature is 1 if the concerned bag has instances falling into the i^{th} group and 0 otherwise. This is essentially a simplified version of our ITE method encoded by training instances only.

As for the SVM predictor in Eq (6), the linear kernel is used for Mi-FV and MaxP due to the large number of samples/dimensions, and the Chi-square kernel² is used for others. A binary two-class SVM model is trained for each emotion class separately.

The experimental results are shown in Figure 4, which clearly demonstrate that our ITE method significantly outperforms the four alternatives with large margins on all three datasets. This validates the effectiveness of our method in generating a better video-level feature representation based on the auxiliary images. In particular, the improvement of ITE over CCE and Mi-FV verifies that the knowledge transferred from the auxiliary emotional image dataset is probably more critical than that existing in the training video frames. This supports our argument that most of the frames of these videos have no direct relation to the emotions expressed by the videos, and underscores the importance of knowledge transfer.

One should notice that CCE has the worst performance. CCE re-encodes the multi-instances into *binary* representations by ensemble clustering. Such representations may have better performance than the hand-crafted features used in [78], but they cannot beat the recently proposed deep

features, which have been shown to be able to extract higher level information [40]. In other words, the re-encoding process of CCE loses discriminative information gained from the deep features, and is therefore unsuited for the task.

In addition, Mi-FV and MaxP have similar performance: MaxP is slightly better on Plutchik, Ekman and Mi-FV is slightly better on YouTube. However, the results of Mi-FV and MaxP are much worse than those of AvgP. These differences can be explained by the different choices of kernels. We validate that the AvgP with linear SVM classifier has similar performance (with a variance of 2%) as MaxP and Mi-FV. Nevertheless, due to high dimensions of Fisher Vectors and large amount of training instances in MaxP, nonlinear kernels will introduce prohibitive computational cost. Thus, in subsequent experiments, we use AvgP as the major alternative baseline to ITE since other alternatives do not demonstrate competitive advantages.

4.2.2 Zero-Shot Recognition

Experiments on zero-shot emotion recognition are conducted on two datasets: the Plutchik and the YouTube. We use *anger*, *joy*, *surprise*, and *terror* as testing classes in the Plutchik dataset (300 testing instances in total). We validate our methods on both versions of YouTube dataset: YouTube-8 and YouTube-24. For YouTube-8, we use *fear* and *sadness* as the testing classes. For YouTube-24, we randomly split the 24 classes into 18 training and 6 testing classes with 5-round repeated experiments. In this zero-shot setting, no testing classes are seen during training.

First and foremost, we need to highlight that without the heterogeneous knowledge transferred from text, there is no way to enable zero-shot emotion learning. There is no previous work that had ever reported the successful experiments on zero-shot emotion learning, mainly due to the difficulties in clearly defining the semantic word/attribute vectors of emotion classes. In this paper, a language model trained from large-scale text data is employed to transfer the knowledge of emotion words to the video domain via the semantic word vectors. Such a transferring step enables the zero-shot emotion recognition.

We compare our T1S algorithm with Direct Attribution Prediction (DAP) [41], [42]. For DAP, at test time each dimension of the word vectors of each test sample is predicted, from which the test class labels are inferred. DAP can be taken as directly using Eq (10) without the word vector smoothing by Eq (9). Four variants are compared: (a) using different video-level feature representation (AvgP or ITE); (b) using different zero-shot learning algorithm (T1S or DAP). Figure 5 shows the results. Our ITE+T1S approach is the best performing method, outperforming the second best by 3.6, 4.8, and 1.2 absolute percentage points respectively. The results confirm the effectiveness of the knowledge transfer scheme, which improves upon the random baseline by 8.1, 6.3, and 15.9 absolute percentage points. The second best model for Plutchik and YouTube-24 is AvgP+T1S, whose performance is close to the best method. This suggests the T1S technique contributes the biggest performance gain when the training classes bear some similarity to the unseen testing classes. However, when the training classes are very different from the testing classes, the ITE encoding scheme plays an important role. Overall, the experiments show the

2. The RBF kernel is also evaluated but shows slightly lower performance than that of the Chi-square kernel.




Fig. 4: **Supervised Learning on emotion dataset analysis.** We use 2000 bins for our method. The two baselines are used both linear kernel/chi-square kernel.




Fig. 5: **Zero-shot Learning results of Youtube dataset.** Illustrated is the Mean Accuracy.

combination of ITE+T1S is effective under different zero-shot learning conditions. Given the inherent difficulties of the zero-shot learning task, we consider the results to be very promising.

4.2.3 Key Implementation Choices

In this subsection, we discuss the settings of two important experimental options in our framework.

Number of clusters for auxiliary emotional images. We vary the number of clusters (i.e., D) of the auxiliary images in Eq (2). These experiments are conducted on the Ekman dataset. The results are plotted in Figure 6, which further validate the effectiveness of our ITE method: when varying the number of clusters from 100 – 5000, our ITE results are gradually improved. This is due to the fact that the increased number of clusters help capture additional discriminative information.

Fine-tuning of the AlexNet CNN model. We also validate the experiments by using the auxiliary images to fine-tune the weights in the AlexNet CNN model. After the fine-tuning step, the classification results on all the three datasets have only small changes ($\pm 0.5\%$) rather than significant




Fig. 6: Results on the Ekman dataset, using different numbers of auxiliary image clusters.

improvements. We postulate that the videos in our datasets have more diversified contents and relatively lower quality than the Flickr images. In other words, there are some differences between the distributions of the images and the videos, which is the reason that fine-tuning with images is not very helpful.




Fig. 7: Quantitative evaluation results of video emotion attribution.




Fig. 8: An example of video emotion attribution.

4.3 Video Emotion Attribution

As discussed earlier, another advantage of our encoding scheme is that we can identify the video clips³ that have high impact on the overall video emotion.

As the first work on video emotion attribution, we define the evaluation protocol of user study to evaluate the performance of different algorithms for this task: Ten subjects, unaware of project goals, were invited for the user study. Given all emotion keywords of the corresponding dataset and clip computed from the video, participants are asked to guess the name of the emotion expressed in the clip. These clips are selected by different methods discussed below. Since the ground-truth video emotion labels are known, we are able to compute the fraction of participants who assigned the correct emotion label for each clip.

We randomly select 20 videos from each of the three datasets. For each video, we compute the 2-second video clip that contains the highest attribution towards video emotion, using Eq (12).

We compare several different methods: **Chance**. It is the prior probability of guessing the emotion without viewing any portion of the video **Random sampling**. We first randomly sample 2 other clips (2 seconds each) from the same video that do not overlap with the first clip. **Face present**. "face_present" features [79] is used to rank all the videos frames. Larger and more faces are detected in those frames that are highly ranked. 2-clips are thus generated by using the top rank frames.

3. Note that in term of our pilot study, the emotions are sparsely expressed in the video. On averagely around 10% of totally video frames are related to emotion on our three dataset.




Fig. 9: User-study results of the video summarization methods.

The results are shown in Figure 7. Our method achieve best accuracy, and outperforms "face_present" baseline by 16–26 absolute percentage points. It is true that facial expression can convey strong emotions. Nevertheless, a large portion of videos in our dataset are user-generated videos which are more general than face videos used in traditional facial emotion recognition tasks. Those videos do not even have the human faces present. This result indicates that our method can consistently identify video clips that convey emotions recognizably similar to the emotion conveyed by the original video. That is, the identified clips contribute to the overall video emotion.

A qualitative result of emotion attribution is shown in Figure 8, where the video is uniformly sampled every 10 frames. The bar chart shows scores of different frames, where the key frames are shown above the bars. The figure demonstrates that clips with stronger emotional contents are given higher scores of attribution, validating the effectiveness of our method.

4.4 Emotion-Oriented Video Summarization

Finally, we evaluate our framework on emotion-oriented video summarization. We compare with four baselines: (1) **Uniform sampling**. It uniformly samples several clips from video. (2) **K-means sampling**. It simply clusters the clips and selects a clip closest to each cluster centroid. (3) **Story-driven summarization** [49]. This approach was developed to summarize very long egocentric videos. We slightly modify the implementation and make the length of the summary controllable for our task. (4) **Real-time summarization** [72]. Wang *et al.* [72] aim at efficient summarization of videos based on semantic content recognition results. For all the methods, the length of summary is fixed to 6 seconds if the original video is longer than 1 minute. For short videos, the length is fixed at 10% of original video duration.

Following [66], we conduct a user study to evaluate different summarization methods. Ten subjects unfamiliar with the project participated in the study. We show the summary results of all the methods (without the audio information) to each participant. Participants are asked to rate each result on a five-point scale for each of the following evaluation metrics: (1) *Accuracy*: the summary accurately describes the "dominating" high-level semantics of the original video; (2) *Coverage*: the summary covers as much visual content using



Fig. 10: The qualitative results of emotion-oriented video summarization.

as few frames as possible. (3) *Quality*: the overall subjective quality of the summary; (4) *Emotion*: the summary conveys the same main emotion as the original video.

The results are shown in Figure 9. The average score is shown in the “Overall” column. Our method (“EmotionVS”) performs better than the other methods on the accuracy and the emotion metrics. On the emotion metric, we beat the best baseline by a margin of 0.87. Although we are doing slightly worse on the coverage metric (-0.13 compared to the best baseline), the drop in quality is minimal (-0.04 compared to the best baseline). The results suggest that the selection of emotional key frames and clips does not only capture the emotion of the original video, but also improves the overall accuracy of the summary, since an accurate summary should capture emotional content. Our emotion-oriented summarization method significantly increases the amount of emotional contents captured by the summary without material loss on other quality measures.

A qualitative evaluation is shown in Figure 10. At the

top, the figure shows a video of an art therapist (the woman in green). Different from other methods, our summarization not only captured the therapy procedure, but also focused on the sadness of the therapist, which is the central emotion conveyed in this video. At the bottom, we illustrate a user-generated video where a father surprises his daughter during a baseball game by dressing as the catcher and revealing himself. All baseline methods are more focused on the baseball game itself, which, however, is the least related to the emotion of this video. In contrast, our method clearly captures the revealing of the father, the surprised daughter, and the subsequent, emotional hug.

5 CONCLUSIONS

This paper provides the first study of knowledge transfer from heterogeneous sources for the task of video emotion understanding, which includes supervised and zero-shot emotion recognition, emotion attribution and emotion-oriented summarization. For effective knowledge transfer,

we learn encoding schemes from a large-scale emotional image data set and a large, 7-billion-word text corpus. This transfer facilitates the creation of a representation conducive to the tasks of video emotion understanding. In zero-shot emotion recognition, an unknown emotional word is related to known emotion classes through the use of a distributed representation in order to identify emotions unseen during training. Our experiments on three challenging datasets clearly demonstrate the benefit of utilizing external knowledge. Our framework also enables novel applications such as emotion attribution and emotion-oriented video summarization. A user study shows that our summaries accurately capture emotional content consistent with the overall emotion of the original video. Future work will address the joint application of emotion-oriented summarization and story-driven summarization, which should allow us to create complete and emotionally compelling stories.

ACKNOWLEDGMENTS

The authors would like to thank Chong-Wah Ngo for his constructive advise.

REFERENCES

- [1] J. Amores. Multiple instance classification: Review, taxonomy and comparative study. *Artif. Intell.*, 201(4):81–105, 2013.
- [2] S. Andrews, I. Tsochantaridis, and T. Hofmann. Support vector machines for multiple-instance learning. In *NIPS*, pages 561–568, 2003.
- [3] L. F. Barrett. Are emotions natural kinds? *Perspectives on Psychological Science*, 1(1):28–58, 2006.
- [4] L. F. Barrett, K. A. Lindquist, and M. Gendron. Language as context for the perception of emotion. *Trends in cognitive sciences*, 11:327–332, 2007.
- [5] D. Borth, R. Ji, T. Chen, T. M. Breuel, and S.-F. Chang. Large-scale visual sentiment ontology and detectors using adjective noun pairs. In *ACM MM*, 2013.
- [6] N. Bosch, S. D’Mello, R. Baker, J. Ocumpaugh, V. Shute, M. Ventura, L. Wang, and W. Zhao. Automatic detection of learning-centered affective states in the wild. In *the 2015 International Conference on Intelligent User Interfaces*, 2015.
- [7] J. M. Carroll and J. A. Russell. Do facial expressions signal specific emotions? Judging emotion from the face in context. *Journal of Personality and Social Psychology*, 70(2):205–218, 1996.
- [8] T. Chen, D. Borth, Darrell, and S.-F. Chang. DeepSentibank: Visual sentiment concept classification with deep convolutional neural networks. *CoRR*, 2014.
- [9] Y. Chen, J. Bi, and J. Z. Wang. Miles: Multiple-instance learning via embedded instance selection. *IEEE TPAMI*, 28(1):1931–1947, 2006.
- [10] A. C. Cruz, B. Bhanu, and N. S. Vision and attention theory based sampling for continuous facial emotion recognition. *IEEE TAC*, 5(4):418–431, 2014.
- [11] D. DeMenthon, V. Kobla, and D. Doermann. Video summarization by curve simplification. In *ACM MM*, 1998.
- [12] A. Dhall, R. Goecke, J. Joshi, M. Wagner, and T. Gedeon. Emotion recognition in the wild challenge 2013. In *ACM ICMI*, 2013.
- [13] A. Dhall and G. Roland. Group expression intensity estimation in videos via gaussian processes. In *ICPR*, 2012.
- [14] T. G. Dietterich, R. H. Lathrop, and T. Lozano-Pérez. Solving the multiple instance problem with axis-parallel rectangles. *Artif. Intell.*, 89(96):31–7, 1997.
- [15] R. J. Dolan. Emotion, cognition, and behavior. *Science*, 298:1191–1194, 2002.
- [16] P. Ekman. Universals and cultural differences in facial expressions of emotion. In *Nebraska Symposium on Motivation*, 1972.
- [17] P. Ekman. An argument for basic emotions. *Cognition & emotion*, 6(3-4):169–200, 1992.
- [18] R. el Kaliouby and P. Robinson. Real-time inference of complex mental states from facial expressions and head gestures. In *CVPR*, 2004.
- [19] Y. Fu. Multi-view metric learning for multi-view video summarization. *CoRR*, 2014.
- [20] Y. Fu, Y. Guo, Y. Zhu, F. Liu, C. Song, and Z.-H. Zhou. Multi-view video summarization. *IEEE TMM*, 12(7):717–729, 2010.
- [21] Y. Fu, T. M. Hospedales, T. Xiang, and S. Gong. Learning multi-modal latent attributes. *IEEE TPAMI*, 2013.
- [22] Y. Fu, T. M. Hospedales, T. Xiang, and S. Gong. Transductive multi-view zero-shot learning. *IEEE TPAMI*, to appear.
- [23] Y. Fu, T. M. Hospedales, T. Xiang, S. Gong, and Y. Yao. Interest-ness prediction by robust learning to rank. In *ECCV*, 2014.
- [24] Y. Fu, T. M. Hospedales, T. Xiang, J. Xiong, S. Gong, Y. Wang, and Y. Yao. Robust subjective visual property prediction from crowdsourced pairwise labels. *IEEE TPAMI*, to appear.
- [25] T. Gärtner, P. A. Flach, A. Kowalczyk, and A. J. Smola. Multi-instance kernels. In *In Proc. 19th International Conf. on Machine Learning*, pages 179–186. Morgan Kaufmann, 2002.
- [26] X. Glorot, A. Bordes, and Y. Bengio. Domain adaptation for large-scale sentiment classification: A deep learning approach. In *ICML*, 2011.
- [27] J. J. Gross. Emotion regulation: Affective, cognitive, and social consequences. *Psychophysiology*, 39(3):281–291, 2002.
- [28] A. Habibian, T. Mensink, and C. G. M. Snoek. Videostory: A new multimedia embedding for few-example recognition and translation of events. In *ACM MM*, 2014.
- [29] A. Hanjalic and H. Zhang. An integrated scheme for automated video abstraction based on unsupervised cluster-validity analysis. *IEEE TCSVT*, 1999.
- [30] G. Irie, T. Satou, A. Kojima, T. Yamasaki, and K. Aizawa. Affective audio-visual words and latent topic driving model for realizing movie affective scene classification. *IEEE TMM*, 12(6):523–535, Oct 2010.
- [31] Q. Ji, P. Lan, and C. Looney. A probabilistic framework for modeling and real-time monitoring human fatigue. *Systems, Man and Cybernetics, Part A: Systems and Humans*, *IEEE Transactions on*, 36(5):862–875, Sept 2006.
- [32] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. *CoRR*, 2014.
- [33] Y.-G. Jiang, B. Xu, and X. Xue. Predicting emotions in user-generated videos. In *AAAI*, 2014.
- [34] Y.-G. Jiang, J. Yang, C.-W. Ngo, and A. G. Hauptmann. Representations of keypoint-based semantic concept detection: A comprehensive study. *IEEE TMM*, 12(1):42–53, 2010.
- [35] Y.-G. Jiang, YanranWang, R. Feng, X. Xue, Y. Zheng, and H. Yang. Understanding and predicting interestingness of videos. In *AAAI*, 2013.
- [36] B. Jou, S. Bhattacharya, and S.-F. Chang. Predicting viewer perceived emotions in animated gifs. In *ACM MM*, 2014.
- [37] H.-B. Kang. Affective content detection using HMMs. In *ACM MM*, 2003.
- [38] A. Kapoor, W. Bursleson, and R. W. Picard. Automatic prediction of frustration. *Int. J. Hum.-Comput. Stud.*, 65(8):724–736, 2007.
- [39] D. Kotzias, M. Denil, P. Blunsom, and N. de Freitas. Deep multi-instance transfer learning. *CoRR*, 2014.
- [40] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- [41] C. H. Lampert, H. Nickisch, and S. Harmeling. Learning to detect unseen object classes by between-class attribute transfer. In *CVPR*, 2009.
- [42] C. H. Lampert, H. Nickisch, and S. Harmeling. Attribute-based classification for zero-shot visual object categorization. *IEEE TPAMI*, 36(3):453–465, 2013.
- [43] B. Li. A dynamic and dual-process theory of humor. In *The 3rd Annual Conference on Advances in Cognitive Systems*, 2015.
- [44] K. A. Lindquist, T. D. Wager, H. Kober, E. Bliss-Moreau, and L. F. Barrett. The brain basis of emotion: a meta-analytic review. *Trends in cognitive sciences*, 35(3):121–143, 2012.
- [45] D. Liu, G. Hua, and T. Chen. A hierarchical visual model for video object summarization. *IEEE TPAMI*, 32(12):2178–2190, 2009.
- [46] G. Liu, J. Wu, and Z. Zhou. Key instance detection in multi-instance learning. In S. C. H. Hoi and W. L. Buntine, editors, *ACML*, 2012.
- [47] M. Liu, S. Shan, R. Wang, , and X. Chen. Learning expressionlets on spatio-temporal manifold for dynamic facial expression recognition. In *IEEE CVPR*, 2014.

- [48] X. Lu, P. Suryanarayan, R. B. Adams, J. Li, M. G. Newman, and J. Z. Wang. On shape and the computability of emotions. In *ACM MM*, 2012.
- [49] Z. Lu and K. Grauman. Story-driven summarization for egocentric video. In *CVPR*, 2013.
- [50] Y.-F. Ma, L. Lu, H.-J. Zhang, and M. Li. A user attention model for video summarization. In *ACM MM*, 2002.
- [51] J. Machajdik and A. Hanbury. Affective image classification using features inspired by psychology and art theory. In *ACM MM*, 2010.
- [52] S. Marsella and J. Gratch. EMA: A process model of appraisal dynamics. *Journal of Cognitive Systems Research*, 10(1):70–90, 2009.
- [53] D. McDuff, R. E. Kaliouby, J. F. Cohn, and R. Picard. Predicting ad liking and purchase intent: Large-scale analysis of facial responses to ads. *IEEE TAC*, 2015.
- [54] T. Mikolov, I. Sutskever, K. Chen, G. Corrado, and J. Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, 2013.
- [55] C.-W. Ngo, Y.-F. Ma, and H.-J. Zhang. Video summarization and scene detection by graph modeling. *IEEE TCSVT*, 15(2):296–305, 2005.
- [56] A. Ortony, G. Clore, and A. Collins. *The Cognitive Structure of Emotions*. Cambridge University Press, 1988.
- [57] A. Ortony and T. J. Turner. What’s basic about basic emotions? *Psychological Review*, 97:315–331, 1990.
- [58] M. Pantic. Machine analysis of facial behaviour: Naturalistic and dynamic behaviour. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1535):3505–3513, 2009.
- [59] R. Plutchik. Emotion: Theory, research, and experience. In *Theories of Emotion*, volume 1. Academic Press, 1980.
- [60] B. Schuller, M. F. Valstar, F. Eyben, G. McKeown, R. Cowie, and M. Pantic. Avec 2011—The first international audio/visual emotion challenge. In *ICACII*, 2011.
- [61] R. G. Schuller B. Recognising interest in conversational speech-comparing bag of frames and supra-segmental features. In *INTERSPEECH*, 2009.
- [62] T. Senechal, J. Turcot, and R. E. Kaliouby. Smile or smirk? automatic detection of spontaneous asymmetric smiles to understand viewer experience. In *IEEE FG*, 2013.
- [63] B. M. Sikka K, Dhall A. Weakly supervised pain localization using multiple instance learning. In *IEEE FG*, 2013.
- [64] J. Sivic and A. Zisserman. Video google: a text retrieval approach to object matching in videos. In *ICCV*, pages 1470–1477, 2003.
- [65] Y. Tian, T. Kanade, and J. Cohn. Facial expression recognition. In S. Z. Li and A. K. Jain, editors, *Handbook of Face Recognition*, pages 487–519. Springer London, 2011.
- [66] B. T. Truong and S. Venkatesh. Video abstraction: A systematic review and classification. *ACM TOMM*, 3(1):79–82, 2007.
- [67] M. Valstar, B. Jiang, M. Mehu, M. Pantic, and S. Klaus. The first facial expression recognition and analysis challenge. In *IEEE FG*, 2011.
- [68] H.-L. Wang and L.-F. Cheong. Affective understanding in film. *IEEE TCSVT*, 16(6):689–704, 2006.
- [69] J. Wang and J.-D. Zucker. Solving the multiple-instance problem: A lazy learning approach. In *ICML, ICML ’00*, pages 1119–1126, 2000.
- [70] M. Wang, R. Hong, G. Li, Z.-J. Zha, S. Yan, and T.-S. Chua. Event driven web video summarization by tag localization and key-shot identification. *IEEE TMM*, 14(4):975–985, 2012.
- [71] S. Wang and Q. Ji. Video affective content analysis: a survey of state of the art methods. *IEEE TAC*, PP(99):1–1, 2015.
- [72] X. Wang, Y. Jiang, Z. Chai, Z. Gu, X. Du, and D. Wang. Real-time summarization of user-generated videos based on semantic recognition. In *ACM MM*, 2014.
- [73] X.-S. Wei, J. Wu, and Z.-H. Zhou. Scalable multi-instance learning. In *ICDM*, 2014.
- [74] C. Xu, S. Cetintas, K.-C. Lee, and L.-J. Li. Visual sentiment prediction with deep convolutional neural networks. *CoRR*, 2014.
- [75] X. Xu and E. Frank. Logistic regression and boosting for labeled bags of instances. In *8th Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 2004.
- [76] Z. Xu, Y. Yang, and A. G. Hauptmann. A discriminative CNN video representation for event detection. *CoRR*, 2014.
- [77] Q. You, J. Luo, H. Jin, and J. Yang. Robust image sentiment analysis using progressively trained and domain transferred deep networks. In *AAAI*, 2015.
- [78] Z.-H. Zhou and M.-L. Zhang. Solving multi-instance problems with classifier ensemble based on constructive clustering. *Knowledge and Information Systems*, 11(2):155–170, 2007.
- [79] R. D. Zhu X. Face detection, pose estimation, and landmark localization in the wild. In *CVPR*, 2012.